

The Mars Climate Sounder In-Flight Positioning Anomaly

Bruno M. Jau* and David Kass**

Abstract

The paper discusses the Mars Climate Sounder (MCS) instrument's in-flight positioning errors and presents background material about it. A short overview of the instrument's science objectives and data acquisition techniques is provided. The brief mechanical description familiarizes the reader with the MCS instrument. Several key items of the flight qualification program, which had a rigorous joint drive test program but some limitations in overall system testing, are discussed. Implications this might have had for the flight anomaly, which began after several months of flawless space operation, are mentioned. The detection, interpretation, and instrument response to the errors is discussed. The anomaly prompted engineering reviews, renewed ground, and some in-flight testing. A summary of these events, including a timeline, is included. Several items of concern were uncovered during the anomaly investigation, the root cause, however, was never found. The instrument is now used with two operational constraints that work around the anomaly. It continues science gathering at an only slightly diminished pace that will yield approximately 90% of the originally intended science.

Introduction

The MCS instrument is one of six scientific instruments flying aboard the Mars Reconnaissance Orbiter (MRO) spacecraft, launched Aug. 12, 2005. The instrument's purpose is to acquire data about the Martian weather and climate. Science observations started Sept. 24, 2006. The duration of mapping operations is scheduled for two earth years but it is hoped that MCS can continue to acquire data for an additional two years. Current information about the MRO mission and this instrument can be found in JPL's [1] and the Planetary Society's [2, 3] web pages. A nearly identical instrument, performing surface science investigations around the Moon, will be launched onboard the Lunar Reconnaissance Orbiter Spacecraft in October of 2008.

On Dec. 11, 2006, about a day after a solar flare near Mars, the instrument's elevation drive started to miss its intended positioning by a few steps several times. The flight software reacts by automatically executing a saving routine after several positioning errors, stowing the instrument in a sun-safe position. Initial evaluations revealed no obvious cause; the position errors were random and happened a few minutes or many days apart. Up to now, over 200 positioning errors have occurred; all with the elevation drive. Scan image evaluations confirmed that most errors were four steps away from the commanded position. A tiger team was tasked to review all software, electronics, and mechanical aspects, to hopefully find the cause of the anomaly and to suggest remedies. This paper gives an account of the MCS in-flight anomaly and its investigations, provides technical background material about it, and lists possible causes. Lessons learned are included for future space instrument builders. The paper should also be of interest for the general public that likes to learn about the intricacy of space science.

Science Objectives and Data Acquisition Techniques

MCS is an infrared (IR) band radiometer, primarily performing atmospheric measurements by measuring the radiance profile versus altitude at the limb (or horizon). One minor component of the science objectives also involves measurements of the surface for thermal and radiative balance calculations [4]. The radiances are inverted via a retrieval process to obtain profiles of temperature, dust, clouds, and

* Robotics and Mechanisms Group, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA

** Earth and Planetary Atmospheres Group, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA

humidity versus altitude. The individual profiles are then combined in 4-dimensional “maps” (covering space and time for each measured quantity) to examine the weather and climate of Mars.

The resulting data will be examined and compared to Martian atmospheric models, much like terrestrial meteorological data. This will provide an understanding of the atmospheric temperatures and winds as well as the ways in which the solar heating directly and indirectly forces them. It will also provide valuable knowledge of Martian climate changes from season to season.

The original plan also called for routine surface observations. This would have represented a small portion of the overall science (~10%) from the instrument. They were primarily for studying thermal and radiative changes of the seasonal polar caps to determine their influence on the weather and climate. The surface observations are also used for establishing the lowest 10 km of the atmospheric profiles. An anomaly-imposed operational restriction currently limits the view to only 10° below the limb, which yields limited, but useful measurements of the surface. Due to the curvature of the surface, this corresponds to an emission angle of 65°, which is useful for some surface science as well as to help constrain the lowest part of the atmospheric profiles.

Detector Configuration

For each of the nine channels (8 IR and 1 visible), the instrument has a 21 element detector array on the focal plane. These are oriented so that they provide a profile of the radiation from the surface to ~80 km when pointed at the limb. Each individual detector is 0.2° by 0.35°, with the narrow dimension providing 5 km vertical resolution on the limb to meet the desired vertical resolution for the profiles.

The radiance profile seen at the limb is a combination of two main factors. The first factor is the amount of emission/absorption; and whether it is opaque at the wavelength of the channel (in which case the detectors do not see all the way through the atmosphere). The second factor is the temperature of the atmosphere. Both tend to decrease with altitude (ignoring the special case of clouds). Combined, these changes tend to produce a radiance profile that stays constant until the atmosphere starts to become transparent, then decreases more or less exponentially (see Fig. 8 for an example profile).

Since the detector signals are not chopped, frequent calibration views of space and the blackbody (or solar) targets are necessary to provide a good calibration of the observations. On the other hand, time spent on calibration views (and the time spent moving between calibration views) is time when no science data is collected, leading to a desire to minimize the calibration observations.

Individual observations are performed with 2 second integration times, but are usually combined to improve the signal to noise in the most sensitive cases. When staring at the limb, the signal to noise and rapid change in signal with altitude allows shifts of as little as 5% of a detector's width to be noticed. The pointing accuracy requirement for the actuators is not that strict since other sources of pointing errors (in particular, spacecraft orientation) also interfere. However, since the pointing can be determined from the radiance profile, the pointing accuracy requirements on the actuator are looser than this.

Instrument Scanning Patterns

MCS slewing is driven by the observations needed to support the science goals of the instrument, including the calibration views. Both the azimuth and elevation actuators are used to achieve the necessary orientations, although the elevation actuator performs the majority of the slewing.

The scanning is described in detail in [4], but a few locations played a key role in understanding the anomaly. A graphical representation of the range and frequency of scanning is shown in Fig. 1. The first position is the stow position. This is with both actuators pointing at 0°. This is a sun safe position that is used when the instrument is powered off, since even a momentary exposure of the telescopes will destroy the focal plane. It is also used when the instrument is standing down from normal operations.

When the elevation is at 0°, the telescopes are staring at the built-in black body, one of three calibration views. The second calibration target, the solar target, is at an elevation of 37°. The standard MCS scan observes the blackbody about five times per hour and the solar target once an hour.

Limb, Space and “Mars Nadir”
drift by $\sim 3^\circ$ over each orbit.

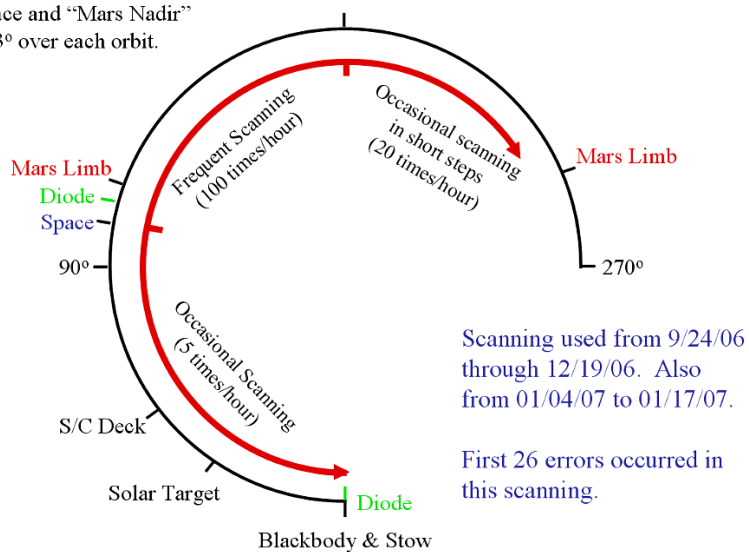


Figure 1: Range of Elevation Actuator Operation

Most of the time, MCS is pointed at the limb (or horizon) of Mars to acquire the key observations of the Martian atmosphere. The location of the limb varies by $\sim 2^\circ$ due to the elliptical nature of the MRO orbit, but is centered on the elevation actuator position of 111.2° . Tied to the limb location is the space view (the third calibration point). It is 8.4° above the limb (or a nominal elevation angle of 102.8°).

When the elevation actuator points between the limb and 250° , MCS is viewing the surface of Mars. In particular, at an angle of 180° , it is pointing in the nadir direction, the sub-spacecraft location.

Scanning is controlled by a series of tables (called SST—Scan Sequence Tables) that contain a list of commands, each indicating the slew and observation duration. These are executed serially by the flight software. The most common scanning is a 35-second cycle that points at nadir for 4 seconds, moves to point at the limb for 16 seconds, then points at space for another 4 seconds before repeating. This pattern was used to determine the lifetime requirement of 1.75M scan cycles.

A second regularly used pattern is the blackbody scan. This scan starts at the space view (usually at the end of a limb scan), moves to view the black body, returns to view space and finishes by slewing to nadir for the start of the next limb sequence.

The high-resolution camera (HiRISE instrument on MRO) placed requirements on MCS, limiting the MCS contribution to the spacecraft jitter that would interfere/blur the HiRISE images. In addition to the jitter requirements, a freeze command was implemented that would point MCS at a sun safe position on the limb during off-nadir HiRISE imaging.

Mechanical Description of the MCS Instrument

Key Design Requirements

The critical design requirements pertaining to the anomaly were:

- *Joint Life: The elevation drive shall survive 3.5M cycles:* This is double the nominal mission lifetime of 1.75M cycles to provide the appropriate margin (and to provide lifetime for an extended mission). This is a significant driver for the design of the transmission and the selection of the lubrication. Pennzane 2000 [5] was chosen as lubricant because of its excellent life characteristics. The life test of a joint proved that transmission and lube could survive the required cycles without much degradation (see [6] and “Qualification Testing” below). Due to scanning limitations imposed since the actuator anomaly, the actual scanning cycles for the primary mission will be significantly lower.

- *Pointing Accuracy: The drives must be able to position the instrument to within $\pm 0.025^\circ$:* The step rate was chosen to be .1 degree; the positioning thus has to be within $\frac{1}{4}$ step, which is well within a stepper actuator’s capabilities. During acceptance testing, all measured pointing accuracies were within $\pm 0.015^\circ$ for both drives over the entire range of motion. This level of accuracy is close to the pointing detection threshold of the detectors.

- *Redundant Position Sensing:* The joint shall incorporate two redundant diodes as position sensors. Diodes are incorporated in the joint drives at 0° and 105.7° . The diodes are regularly checked in-flight. Neither has ever malfunctioned.

- *Position Holding:* The drives shall hold their position when power is removed: The actuator's detent torque, providing 18in-lbs (2Nm) of holding torque, served as launch lock and holds the joint's position during observations. No drifting has ever been detected during sensing.

Mechanical Construction

The following is a brief mechanical description of the instrument. A more detailed description of MCS can be found in publication [6]. The completed instrument is shown in Fig. 2. As can be seen, the entire instrument is covered by thermal blankets.



Figure 2: The Completed MCS Instrument





Figure 3: Instrument without Blankets

Fig. 3 shows the instrument without its thermal blankets. It consists of an optical bench assembly, two rotational drives that provide azimuth and elevation scanning capabilities, and a yoke-like structure that serves as structural support for the upper part of the instrument. All motor control electronics resides on three individual electronics boards located inside the yoke structure.

The optical bench is a thin walled rectangular structure, containing two telescope assemblies. One of the telescope assemblies is depicted in Fig. 4. The focal plane is at the inside surface of the hexagonal block at the upper right. A cross sectional view through the optical bench is shown in Fig. 5: The incoming light passes through the baffle first. The beam is then relayed by a 3-mirror optical beam system before reaching the detector arrays at the focal plane. The telescope's control electronics is on two electronics boards, one attached to the top, the other to the bottom of the optical bench.



The solar calibration target is the bright plane that can be seen in Fig. 3. It reflects indirect sunlight into the instrument. The black body calibration target is recessed in the yoke structure. The instrument can calibrate itself against the black body calibration target when the instrument is pointing straight down.

Figure 4: One of two Telescope Assemblies




Figure 5: Cross Sectional view through the Optical Bench

Joint Drives

Both joint drives [7] combine an actuator and a twist cap section in a compact design. Their actuators are identical but the output bearing configurations are different: The azimuth joint has a hub extension so that its bearings could be spaced further apart; and it features dual duplex bearings. This enabled it to withstand the more severe launch loads at the base of the instrument. To further increase the bearing's load capacity, all output bearings feature full complement ball bearings, thus eliminating the usual bearing retainer (see "Keystoning Effects" below). A final difference is that the azimuth joint has the bearings in a back-to-back configuration for stiffness, whereas the elevation joint has them aligned front-to-front, enabling compliance to the third bearing at the other side of the optical bench. Fig. 6 depicts the azimuth joint with its output bearings spaced further apart. An azimuth-type joint was used for the life test.

The motor is a brushless stepper motor, having a step angle of 30 degrees. It has redundant windings and detent capabilities. Its stall torque is .012Nm. The transmission is comprised of a planetary gear stage, followed by a harmonic drive. The overall gear ratio is 297:1.


The twist cap chamber contains six printed circuit flex cables in a clock spring arrangement, enabling the electrical wiring to transition from the stationary to the rotating section. Each flex print has 14 electrical traces for a total of 78 power and signal traces and 6 grounding traces.

Position Sensing

There are two diode assemblies per joint, sensing the joint's position at 0° and at 105.7° degrees. The 0° position is used for re-initialization. The 105° position is between the limb and space views, and is crossed twice every 35 seconds.

A disk (near item 5 in Fig. 6) normally blocks the light path between the light emitting diode (LED) and the phototransistor. The disk has one slot, allowing diode light to pass through when the slot passes a sensor. The slot size had to be made quite large (approx. 1-1/2mm) for sufficient light from the diode to pass through. Therefore, the sensor is triggered for up to six steps during each diode passage.

During initialization, the 0-step is determined by scanning past the 0° diode first while recording all positions the diode is triggered, then computing the center position. By evaluating the diode and knowing which of the four motor phases corresponds to the computed center step, the controller can uniquely determine the joint's position. Subsequent position sensing only checks if the sensor is triggered when the software anticipates passing the previously determined center step of either diode. To save diode life, the diodes are only turned on when the system expects to pass one of the sensors.



If the actual joint position does not correspond to the commanded position (due to whatever causes the anomaly), the position error has to be at least 3 steps away from the commanded position to be detected by this sensor system, having a 6 step diode trigger window. A sensed error reading is recorded but not reported immediately, thus complicating the error analysis (see “Error Discussion” below). Error cancellation is also possible if a subsequent error is in the opposite direction. However, if the flight software detects an error, it will reinitialize (see below), even if a subsequent error (hysteresis effect), should cancel it out.

Figure 6: The Azimuth Actuator/Twist Cap

Motor Control

Micro-stepping motor control was implemented to reduce motor vibrations that otherwise would propagate to other sections of the spacecraft, in particular to the highly sensitive HiRISE camera. The flight motor controller is a 4-phase micro-stepper controller: At low speeds, the motor driver ramps the currents in a quasi-sinusoidal fashion, rather than by discrete square waves. A complete sine wave advances the motor by four cardinal steps. Hence, an electrical cycle is comprised of 4 phases. At higher speeds, the current input per motor step switches to fewer discrete micro steps. The number of micro steps diminishes to two with increasingly shorter step duration when the motor approaches full speed.

For the error analyses, it is important to note that, if the energized motor functions as expected, it is in sync (or in-phase) with the motor controller at nominally zero position errors. Should the motor lose steps, its speed most likely will drop to zero or it will automatically re-synchronize to the controller at a multiple of four steps from the commanded position. Since most of the measured position errors were four step errors, it follows that the system stops one control sequence away from the commanded position. This does not explain, however, the few other position errors for which telemetry indicated another step difference from the expected position, even some confirmed position overshoots. A mechanical cause, producing substantial overshoots, is highly unlikely.

Qualification Testing

Common wisdom says: test what you fly, fly what you test. But reality is not that simple: A full instrument qualification program was undertaken for MCS, as is customary for flight hardware, including a separate life test for one drive joint. Both flight joints were independently acceptance tested as well. The purpose of this section is not to give a full account of the test program, but to highlight some limitations the test program faced, and the consequences this might have had for the anomaly and anomaly investigation.

- *Tests not done in zero-G:* this was universally considered as acceptable. The issues were that the life test had to be done with the joint axis aligned horizontally. Otherwise, the lube might have gravitated out of the bearings and the flex tapes would have sagged and rubbed against the twist cap walls.

- *Life test not done in vacuum:* This remained a contested item throughout the development with ever reoccurring questions at reviews, particularly concerning lube issues and heat dissipation. There were just simply not enough resources to do a full-simulated life test, lasting 18 months, in vacuum: Heat dissipation concerns could be settled by doing the necessary analytical work. The lube issues were partially resolved by scientific evidence, partially by pointing to similar devices that used the same lubricant and successfully operated in space.

- *Joint testing not done with a flight-like controller:* In today's short development cycles, there often is not enough time to develop long lead components in time for integrated testing: A duplicate flight-like controller was not available for the life test, nor was a flight-like controller ready in time for the qual/acceptance testing of the joint drives. For all joint stand-alone testing, an off the shelf, 2-phase pulse with modulated (PWM) stepper motor controller, emitting square wave signals, was used. The flight controller uses micro stepping instead.

It was determined only during integrated instrument testing that the motor produced substantially less torque than anticipated with the micro stepper controller. This jeopardized much of the previous qualification and acceptance testing. Re-testing was not possible on the integrated instrument and renewed joint stand-alone testing was impossible due to time constraints. In particular, the input waveform generated by the flight controller was insufficiently measured, analyzed, or understood.

To compensate for the reduced torque, an in-flight remedy was implemented, increasing the in-flight minimum temperature for the elevation actuator to 15degC. This keeps the viscosity of the actuator's lube lower, offsetting the decrease in torque with the micro stepper controller. Fortunately, enough heaters and additional heating power could be made available. What remained unsettled, however, was that the actual torque the motor produces with the micro stepper driver was not well known, nor was the effect of using different electrical motor inputs sufficiently well understood and evaluated.

One position error occurred during a routine aliveness test in spacecraft assembly and testing. It was unfortunately not detected until a review of the telemetry a few days later. A visual inspection of the instrument revealed no issues. The instrument performed as expected and all other telemetry was as expected. Due to limited resources and time, no further investigation was performed at that time, although instrument scanning was monitored for a repeat occurrence (none occurred prior to launch).

The obvious 'lesson learned' here is to allow enough time for adequate overall system testing and to have spare resources. The follow-on Diviner development proved that this is difficult. Schedule slips and late delivery of components again required substituting flight components with dissimilar engineering model (EM) components, causing the overall end-to-end system testing to be shortened as well.


Joint Drive Operational Parameters				
Parameters	Life Test	Ground Test	Azimuth Drive	Elevation Drive
	MCS	MCS	MCS in Flight	MCS in Flight
Nr of load cycles required	3,500,000	N/A	350,000	3,500,000
Env. or Op. Temperature degC	-2, +3, RT, +60°C	-3	5 to 15°C	15 to 25°C
Load Velocity steps/sec	450	380	263	263
Load Acceleration deg/sec ²	25	25/42	25	42
Load Inertia kgm ²	0.1045	.12 to .14/.03	0.12 to 0.14	0.03
Voltage V	28	28	32.5	32.5
Electrical Input	PWM	micro stepping	micro stepping	micro stepping
Environment	Nitrogen	air	vacuum	vacuum

Table 1: Life Test Joint vs. Space Operations

- *Differences between life test and space environment:* Table 1 lists the differences between the life test and the actual in-flight conditions for all MCS joint drives. In summary, the load inertia was close to the value for the azimuth joint, but the 3.5M load cycles reflect the life of the elevation drive. With the life tested joint surviving 3.5M cycles at a torque level equal to twice the value the elevation drive requires, and by using a square wave controller that produces much larger dynamic load spikes (hammering

effect), it can be stated that the life joint was over-tested. The testing validated the mechanism and flex cable's ability to function properly far beyond the full life requirement of this mission. Moreover, it demonstrated a capability that far surpasses previous NASA flight experiences with similar space flight mechanisms. During the entire life test, the drive stalled only once at 2.3M cycles at -2degC. The cause was increased viscosity of the lube because of a slight degradation of the lube due to aging. Cold testing was resumed at +3degC and finished without any further incidence.

Fig. 7 shows a close up picture of the harmonic drive, taken after the life test: No mechanical wear could be detected at any mechanical parts, nor was there any graceful degradation evident, which would have expressed itself, for instance, by substantial increases in backlash. But the lube's color had changed from brown to black in the engagement zones between gear teeth and in the bearings. Chemical analyses,



however, revealed that lube samples taken from inside the engagement zone had retained their lubricating capabilities. The joint thus maintains its full capabilities.

The overall conclusion based on life test results (for life test conditions), is that the joint demonstrated its capability to survive the required mission error free, pointing away from a mechanical cause for the anomaly: The life test joint is still being used on occasion today, proving its mechanical integrity and the lube's lubrication capabilities time and time again. It runs error free, provided it receives the proper electrical input signal and current.

Figure 7: Close up View of Mechanical Components after the Life Test

Keystoning Effects

During the assembly of Diviner's elevation joint, it was noticed when the partially assembled joint was rotated freely (actuator not attached yet), that the output bearings occasionally experienced some hard spots for no obvious reasons. The bearings ran smoothly again when part of the inner ring preload, applied by the bearing retainer (at arrow tip of item 4 in Fig. 5) was removed. Initially, a non-flat retainer surface was suspected that could have deformed the thin cross section bearings. But even after re-machining the retainer shoulders and checking flatness and perpendicularity of every shoulders of the bearing assembly (all within $\pm .0003$ in, .008mm), hard spots could still be detected occasionally when the full retainer preload of approximately 120lbs (54kg) axial load was applied. This resistance was somewhat position dependent and lasted over a short distance of up to 30 degrees, but could disappear and reappear at different locations.

Not finding another cause, the bearing supplier [8] was contacted. Their engineering department related that they had experienced similar behavior on other applications with full complement bearings as well. The phenomenon is called keystoning. Here is a somewhat personal explanation of this writer for what happens: If the balls are not rolling exactly on the same race diameter, which can happen with angular contact bearings, the rotational velocity of the balls inside the bearing are not exactly the same. Hence, with no retainer to separate the balls, they progress from contacting each other to pressing against each other in the high load zone. That causes sudden frictional forces between balls. This friction force is not sustained over a long distance because the external load on balls fluctuates (i.e. if the bearing's external load is applied radially, then the balls roll away from being under the high radial load zone). After the external load on a ball is removed, it is free to re-align itself in the race grooves and w.r.t. other balls. Thus, a friction load may or might not build up when balls enter the high load zone.

With a retainer present, a ball might temporarily push against the retainer's finger that engages between two balls. But since the retainer is much more compliant, it will not produce excessive friction forces in the high load zones. Thus, the retainer prevents keystoning from occurring. Lesson learned: avoid using full

complement bearings, especially for dithering motions that rotate the bearing just by just a few degrees, even if it means paying the mass penalty for going to a larger size bearing with a retainer.

For Diviner's elevation drive, the remedy was to reduce the pre-load the inner bearing retainer ring exerts. This was possible because this joint does not require the stiffness the azimuth joint needs. The Keystoning behavior noticed during the Diviner assembly was not noticed during the MCS assembly. However without objective evidence that it is not present on the MCS elevation actuator, it is impossible to say that the MCS actuator does not exhibit the same keystoning behavior.


Actuator Pointing Error Detection and Response

Instrument Position Error Detection and Response

Position errors occur when the MCS flight software (FSW) and the hardware disagree on the position of an actuator. The position of the actuators is checked when passing through either of the diode locations. Position errors only occur while slewing, but a mis-pointing caused by a position error will persist (with the same magnitude) until the actuator is re-initialized.

When the MCS FSW determines there is a position error, it sends an error message in the engineering telemetry. After reporting a position error, the FSW re-initializes both actuators, removing any pointing errors. While re-initializing the actuators, slewing is done at a very slow rate and errors are generally irrelevant since the FSW is searching for the diode. After a certain number of position errors, the instrument fault protection stows the instrument and waits for ground intervention.

The complexity is that the FSW does not report the error when it is detected; it is only reported at the end



of an SST (sequence of slews). Thus, the actual pointing error occurred at some point before the last diode crossing in the SST. This also means that if the position error occurred early in a long and complex scan pattern, the instrument will be mis-pointed until the end of the SST.

Limb Validation of Pointing

Without an absolute position encoder, it is not possible to determine the size of a particular error. However, the team has been able to use the instrument data to determine roughly how many steps the instrument is out of position by looking at the radiances when pointed at the limb. This is

done by using channel A3 (core of the 15 micron band) and comparing the calibrated radiances of the limb views during the period around when the error is reported. We use limb view before the probable slewing error, while the error is in effect, and after the actuator has been re-synchronized.

Figure 8: Detection of Slew Errors at the Limb of Mars

Fig. 8 shows an example analysis. For each observation, MCS acquires eight samples of the limb (thus the clusters of 8 lines). The actual detector values are the + symbols, placed on an altitude scale based

on spacecraft (S/C) pointing and assuming the pointing is correct. The red and green families are the two measurements before the error occurred. The light blue profiles are after the pointing error and the brown family is after the actuator was re-initialized. The most accurate determination is done where the radiance is changing the fastest with height. In this case, the 50 to 60 km region was used. It is obvious that the light blue family is mis-pointed. The dotted horizontal line indicates detectors that should be at the same altitude and have the same radiance. The vertical dotted line shows the actual family of detectors that have the same radiance, and are pointed at the same altitude. The two detectors between the actual and expected alignments indicate a pointing error of two detectors. Since each detector is approximately 0.2° tall (or 5 km on the limb), this is a pointing error with the actuator pointing 0.4° (equivalent to 4 steps).

This provides additional information on the timing of many position errors by indicating that the actuator was already mis-pointed while observing the limb, thus reducing the number of candidate scans for the error. The limb view radiances also sometimes (as in Fig. 8) provide knowledge of the magnitude and direction of the error. Based on experience, we have found that for small errors ($< 1^\circ$), looking at the limb radiance provides a pointing accuracy of one actuator step (0.101°).

There are cases (especially over the poles) when the atmosphere is changing too fast with latitude and it is not possible to accurately measure the pointing error. The profiles are too different to determine which detectors are actually pointing at the same location. In many of these cases, it is still fairly obvious that the actuator is mis-pointed. There are also cases where there are no limb views during the scanning sequence before the error is fixed by re-initializing the actuator. In these cases, the size of the position error cannot be determined.

Actuator Anomaly Timeline

Initial Flight Errors

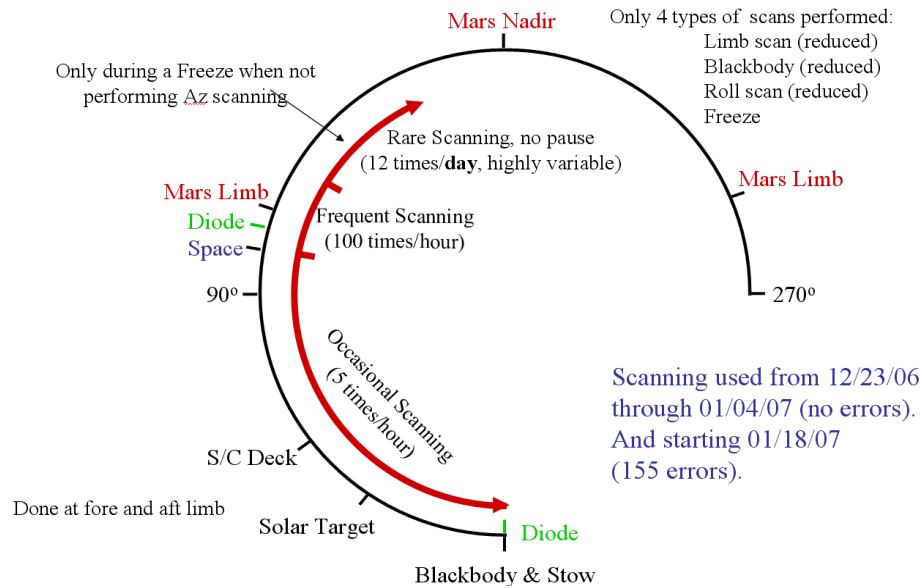
MCS started primary science phase operations on Sept. 24, 2006. After an initial flurry of scanning changes, routine operations settled in. On the evening of Dec. 11, 2006, MCS FSW detected a position error and re-initialized the actuators. Over the next 4 hours, three additional position errors occurred and caused the instrument fault protection to stow the instrument as designed. This occurred at ~5% of the tested life of the actuator.

Over the next two days, an analysis of the telemetry during the position errors showed no other signs of a problem. The initial study of the science observations by the team indicated that there were no large pointing errors in the data. The first four errors were very intermittent and all showed preconditioning by the black body viewing (see below). The last instrument ground commanding prior to the error (a routine health status check) had been 3 days before. The last change to the scanning had been 11 days before. All of the commanding was verified to be correct and not the cause of the position errors.

Initial Response and Diode System Testing

An initial fault tree was defined. The most likely candidate on the tree was the elevation 105° -diode system (since this diode had reported all four errors). The other early candidate was the occurrence of a solar particle event about a day before. Three initial actions were taken: Another health status check was performed (it was unchanged). Second, on Dec. 14, MCS fault protection was initialized with a count of 6 position errors (stow on the 7th error) and scanning resumed. This was to see if the errors were transient. And if not, the scanning would hopefully provide additional information on the actuator health. After 7 hours of scanning, MCS FSW had detected 7 more position errors and again stowed the instrument. This scanning did show that the blackbody preconditioning (see below) was not required for an error to occur. It also indicated that errors could occur in sequences other than the limb scan.

In order to test that the problem was the 105° -diode system, a quick patch was developed to disable its use by the FSW. On Dec. 19, this was ready, installed, and scanning resumed with the position error counter set back to its default value of three. The expectation was that this would stop the FSW from reporting position errors, allowing the investigation to focus on the 105° -diode system while the instrument continued to collect science observations. If the position errors were due to a real mis-pointing of the actuator, then they would soon be caught by the 0-diode and corrected by the re-initialization. After approximately 6 hours of scanning, 4 position errors had been detected by the 0-diode.



Initial Diagnosis Tree and Reduced Scanning

A “problem scan” diagnosis tree was developed to identify the slew or slews causing the errors. It was also designed to determine if there were operational workarounds for the position errors. The first step was to use a reduced set of scanning (Fig. 9) that only involved using a blackbody scan and a limb scan, both limited to 10° below the limb. For sun safety reasons, freeze

commands were still allowed to cause occasional scanning up to 35° below the limb. This scan pattern was designed to rule out one of the three candidate scans. Additional steps were designed to differentiate between the other two scans.

Figure 9: Range of Reduced Scanning

The reduced scanning was started on Dec. 22; it ran over the holidays without producing any errors. The next steps of the tree were started on Jan. 2, but still failed to produce additional position errors. Concerned that the results of the steps in the diagnosis tree were misleading the full scanning pattern was tested on Jan. 4. This did not generate any position errors until the morning of Jan. 16.

Initial Ground Analysis

In the intervening time, a number of analyses and ground tests were performed. Testing involved using the MCS life-test actuator and the EM micro stepper controller to attempt to generate position errors. Unfortunately, this controller deviates from the flight controller in a number of ways, shedding some doubt about the meaningfulness of the test results. The initial effort focused on reduced bus voltages (showing a very sharp transition from scanning without errors to not scanning). Attempts to corrupt the scan control software to introduce position errors met with limited success. Failing to execute a slew step at high speed would produce a 4-step position error. Reviews of the software design indicated that short of very targeted and special corruption, there were no ways to mis-command the actuator. And if such corruption existed, position errors would occur on all scans longer than 16°. Reviews of the electronics revealed that much of the firmware and hardware was shared by the two actuators and none of the separate components had known intermittent failure modes. The mechanical review also revealed no obvious causes or mechanisms for generating position errors. During this time the detailed analysis necessary to detect small position errors on the limb was also developed.

Initial January Response and Frequency Study

The next step on the instrument was to repeat the tests from December to see if they made the errors (temporarily) disappear again. In doing this, an attempt was made to isolate the part of the process that had eliminated the errors. The first new scanning, on Jan. 17 involved overwriting parts of MCS random access memory (RAM) to eliminate any corruption. This did not eliminate position errors. A second step taken on Jan. 18 was to return to the reduced scanning. Again, position errors occurred. In these two tests, the error frequency seemed to have increased from 1/hour of scanning to 2/hour, although the time between errors was still highly variable.

The next experiment was designed to better understand the error statistics and to obtain a larger population of errors to examine for trends and correlations. The previous testing had generated 33 errors. The fault protection counter was increased to 50 and reduced scanning was resumed. MCS scanned for almost 3 days before generating the 51 errors and stowing. For sun safety, the scanning was pointed at the aft limb instead of the standard forward limb. This also changed the thermal environment of the elevation actuator. The time between errors varied from 2 minutes to 13 hours. The average was 1/hour (2/hour if the 13 hour gap was not counted). There was a hint that the frequency was increasing slightly towards the end of the 3 days. This run also had the first occurrence of an error immediately following the re-initialization of the previous error.

Azimuth Scanning

Next, the investigation proceeded to attempt to isolate the error by performing extensive (and identical) scanning in azimuth to see if errors would occur. This would either implicate or exonerate the shared components. The plan for test #8 was to scan for one orbit in elevation as usual and then slew in azimuth using the same scan pattern for the next orbit. Due to the control limitations of the instrument, an occasional slew in elevation would occur during the azimuth orbits to respond to freeze commands.




Figure 10: Error Rate in Late January Testing

The results are shown in Fig. 10, plotting the time of the error versus error number. The yellow crescents are the orbits with azimuth scanning. All of the errors during the azimuth orbits are associated with elevation scans. Part way through, it was decided to proceed with the next test (test #9) and remove the last difference between the two actuators by overwriting the elevation slew table with the azimuth table.

A number of features appeared in the test. First, there were no position errors during azimuth scanning. This implicates the drive electronics not shared by the actuators and the elevation actuator itself. During the analysis, it was discovered that on 5 occasions, a much larger error had occurred (they ranged from 2.6° to over 11°). All five errors were associated with slews going more than 10° below the limb. The direction of all the errors was consistent with either stalling on a slew or many 4-step small errors on the way to 35° below the limb (~145° in elevation space).

The other trend noticed during these tests (shown in Fig. 10) is an increase in the error frequency. And even more disconcerting, the rate of increase appears to accelerate. A careful analysis showed that the

major jump in error rate occurred when the first large position error occurred. Note that even at the most rapid, the average was still over 20 slews between errors.

One interpretation of the two tests (in addition to eliminating a large number of systems from the fault tree) was that the elevation actuator was rapidly degrading and would soon be inoperable. The final error rate was also significantly degrading the ability to perform good science observations. And the large pointing errors introduced a sun safety issue. This led to the formation of an anomaly tiger team. At this time, Starsys [7] was actively brought into the investigation.

Limb Staring and Ground Testing

It was not known if the progressive degradation was due to use or just due to aging. Given the fact that staring at the blackbody provided no science, it was decided to move once (Feb. 9) to stare at the limb and avoid using the elevation actuator until the anomaly investigation could perform a thorough analysis. The theory was that if the actuator were to remain jammed in that orientation, it provided the maximum science return (estimated to be about 70% of the planned science).

One series of key tests was a set of torque measurements and controller waveform sampling. These were performed with the EM controller with either the life test actuator or a test motor. Testing revealed that the micro stepper controller using the flight parameters generated significantly less torque than expected. This system no longer met (and did not come close to meeting) the design requirement of a factor of 2 for torque margin at the lower end of the 26 V to 36 V S/C voltage range. There were a number of other peculiarities in the power waveforms supplied to the actuator, mostly related to limitations of the motor controller processor.

Most of the torque reduction was traced to a miscommunication over the back-EMF (electromechanical force) constant to be used for the flight motor winding configuration. A new parameter was determined by analysis and then refined experimentally with the life test actuator. With an increased back-EMF constant, effectively supplying more current, the motor torques were still marginal, but at least now met design requirements at the nominal operating voltage of 32.5V.

In late March, MCS was power cycled due to an MRO safe mode entry. After the power cycle, MCS returned to limb staring. This removed corruption concerns and a number of other very unlikely items from the fault tree.

The in depth investigation determined there were no plausible catastrophic failures on the fault tree. The investigation also developed fixes for a number of the torque and slewing issues (although not all could be fixed, some were built into the way the MCS FSW controlled the motors). But the investigation was not able to point to a root cause, or even a plausible root cause.

Continued Flight Actuator Testing

It was decided to resume using the elevation actuator in late May, incorporating the changes in the motor control parameters to increase the instrument torque and clean up the waveform to the extent possible. The scanning was further reduced to only involve the blackbody (but only twice per hour), space views and the limb. The first four blackbodies were followed by a limb scan that had a position error, ending the test. Due to the simplicity of the scanning, the limb radiance analysis indicated all four errors appeared to be overshoots (the only other scenarios involve complex inaccuracies of the 105° diode system and scans ending exactly the “right” distance short).

The instrument returned to limb staring at this point. An evaluation of the fault tree in light of overshoots focused on actuator resonances (the only fault on both the overshoot and undershoot branches). The blackbody pre-conditioning of all errors (as opposed to January when errors were so frequent, many were occurring without the pre-conditioning) also led to considerations of slew length. The next test (on June 14) reduced the top speed of the actuators to avoid any resonance effects. The change to the top speed did not appear to have any effect; three errors occurred within 30 minutes and the instrument stowed.

Following the end of the previous test, scanning was resumed, but only between the limb and space (~8° total slew length). This was an attempt to see if eliminating the blackbody pre-conditioning improved the

performance. A first position error occurred after 30 minutes, followed by two more 4 hours later. No errors were seen after the third error. Short scans, while not immune to errors, seemed to have improved the performance by eliminating position errors. This scanning continued for the next week without errors.

On June 22 (after developing a new diagnosis tree), the 8° scanning was extended 20° above the limb. The goal was to see if scans that reached top speed failed. They did not fail and after running for 24 hours, it was decided to try the May 31 scanning to verify that position errors had not disappeared again. Unfortunately, 4 errors occurred within approximately two hours and the instrument was returned to scanning only between the limb and space. Again, there were no position errors in this range of scanning.

To try to pin down whether the errors were position dependent or slew length dependent, the scanning was modified on June 28 to add blackbody sequences, but to do the move to and from the blackbody in a series of 8° slews instead of one long slew. This was not successful with errors still occurring on the limb scan following the blackbody scans (although at about half the rate as earlier in June). This seemed to point to a position dependent error generating mechanism. The instrument was returned to a state where it was scanning between the limb and space, and the instrument remained free of errors.

Errors Disappear after Reduced Scanning

Over the July 4 week, one blackbody per day was added to the scanning. This was inserted due to keystone concerns since an 8° slew does not move the balls of the output bearings by a full diameter, it was essentially just rocking them back and forth. This scanning had been performed for most of the previous two weeks and was expected to continue for at least another week over the holiday. The expectation was that one position error would be generated every day or two and if the frequency increased rapidly, instrument fault protection would cut in and stow the actuators.

No position errors occurred. This led to experimentations (over the next two weeks) where the timing and frequency of the blackbody calibrations was varied (reaching the original 5/hour rate). None generated position errors. It seemed that scanning over a limited range for an extended period with the updated motor control parameters (higher torque, cleaner waveform), had made scanning elsewhere less prone to errors. Or perhaps scanning over any region above the limb made scanning in that region less prone to errors. Yet, scanning in January seemed to have contributed to the increase in error frequency.

Resumption of Calibration Views

In late July, a spacecraft issue led to the actuator anomaly investigation going on hiatus until early October. During this time, scanning between limb and space with occasional blackbody scans (done in short 10° scans due to the spacecraft issue) was performed. Also, further EM testing revealed that the new motor controller parameters were causing the controller to current limit when the spacecraft bus voltage was at 32.5 V. The clipping distorted the current waveforms and had the potential to introduce a new source of position errors. This led to the design of a new set of motor controller parameters with slightly lower peak currents that did not clip on the life test actuator in the lab, but still had a slightly higher torque than the January tables and the cleaner waveform.

Starting in October, a set of steps was performed to return to more normal scanning with the new motor controller parameters. This was done gradually with several days between each change in configuration or slew profile. These steps included increasing the number of blackbody views back to 10 and adding in a view of the solar target. However, the elevation actuator was never allowed to go higher than 120°. Combined, the addition of these observations allowed the science channels to be fully calibrated.

Final Flight Testing—Returning Below the Limb

With full calibration restored, a new round of careful testing was started. Given that the errors seemed to be at least partly position dependent and could possibly be improved by repeated scanning, it was decided to slowly expand the scanning below the limb (>120°). During this process, the fault protection was set on a hair trigger to stow the instrument on the first error. This was the region where large errors had previously occurred. Since large errors seemed to be associated with increased error rates elsewhere, there was a desire to minimize the number that occurred.

The first steps was to return to reduced scanning, but eliminate all possibility of going more than 10° below the limb. This was accomplished by some very careful analysis to insure sun safety. No position errors occurred while scanning 10° below the limb over four days. Therefore, the next step was to expand the range to 20° below the limb. Within 20 minutes (and 30 limb scans), a position error occurred and instrument fault protection stowed the instrument. The error was determined to be a small (4-step) position error.

The next day, the scanning pattern was reduced to go only 10° below the limb. This was successful and no further position errors have occurred as of this paper.

Discussion of Errors and Error Statistics

MCS has experienced 203 position errors. The errors fall into 5 groups according to when they occurred (see above). The first error was prior to launch. The first flight family occurred in Dec. 2006, a second family in Jan. 2007, a third family in June 2007 and the final position error occurred in Oct. 2007.

All of the errors have been characterized to the extent possible, given the limited in-flight telemetry. Unfortunately, an extensive analysis of the instrument and spacecraft telemetry associated with the errors has been mostly unproductive, both on an individual basis and in statistical analyses.

One of the difficulties in the investigation is that the specific slew with the position error is almost never known. There is almost always at least a pair of suspect slews for any given position error. In some cases, there are many possible slews where the error could have occurred. And given the slew pairs are in opposite directions, in most cases, an undershoot on one slew or an overshoot on the other could be responsible for the identified error. The May and June diagnostic tests were set up to minimize this uncertainty. But without an absolute position indicator, not all of the uncertainty could be removed. However, there are a few significant points that the analysis did uncover.

Even though the errors families appear indistinguishable, various error-type subsets stand out:

- The five large errors (errors where pointing is off by more than 1 degree) all occurred in January. All of these large errors are known to have occurred at elevation angles > 120°. Since slewing operations were stopped in Jan. 2007, the instrument was only slewed past 120° for ~20 minutes in October.
- Most of the pointing errors that could be measured are 4-step errors. The preponderance of the 4-step errors is assumed to be due to the 4-phase nature of the MCS controller. At nominal speed, a four-step error could potentially be generated by an approximately 15 m-sec lasting signal disturbance. But there are at least four errors that are 2-step errors (and a small family that may be 5 or 6 step errors). One possible way of generating 2, 5 and 6 step errors is to have the error occur at the end of the slew.
- The time (or scanning) between individual errors is highly variable, ranging from 2 minutes (the first scanning after re-initializing) to over 13 hours in one case. Even within a given time frame, the time between errors is highly variable (for example, the 13-hour interval followed a 2-minute interval). While errors may have occurred on consecutive slews, they never occurred on more than 5% of the slews, and in most periods, the frequency is below 0.1%. The general error frequency increased during the December and January groups of errors. During the May/June diagnostic period, the error frequency decreased until the instrument reached an error free state using the current set of environmental constraints.
- No correlation of errors with any environmental condition has been found (at the resolution of the available data). In particular, there are no significant temperature fluctuations. The only possible exception is the spacecraft bus voltage: For most of the MRO orbit, the voltage is at 32.5 V, but it dips as low as 28 V during the eclipse (slowly dropping at eclipse and then returning to 32.5 V immediately afterwards). There are relatively few errors at low voltages (since the bus is at 32.5 V most of the time). At most, the probability of an error at 28 V is twice that above 32 V, but this might not be statistically significant. The life-test actuator (and the limited testing of the flight actuator), indicate there is a very abrupt transition between completely stalling (at 22.8 V) and running fine (at 22.9 V) with the micro

stepper controller. Thus, any in-flight voltage fluctuation should not cause errors. Due to the rapid increase in torque with applied voltage, this also indicates that there is probably ample torque margin when most errors occur.

- The errors appear to occur with any slew lengths. A rigorous exploration of slew length dependency was not undertaken (ambiguities between slew lengths and error rates makes the existing data difficult to interpret). Errors definitely occurred on both short slews and long slews (short slews are $< 16^\circ$ and do not reach top speed). Several of the four errors definitely occurred on slews of $\sim 8^\circ$. It is possible that errors preferentially occur on long slews, but this probably only indicates that the error probability correlates with the distance covered.

- It appears that at least one error can be associated with a slew in each direction. Also, at least one error is an undershoot. There are four errors that (assuming the diodes can detect all combinations of a 4-step error) are overshoots. But without the exact knowledge when slew errors occurred, it is often not possible to determine the direction of slewing, nor whether the error was an overshoot or undershoot.




Figure 11: Blackbody then Limb Error Pattern

Error "Preconditioning"

One peculiar pattern reoccurs frequently as position error: Many of the position errors are in the limb scan sequence following a blackbody sequence. Figs. 11 & 12 illustrate the pattern in two different ways. Due to crossing the diode in the blackbody sequence, and the fact that there is no error reported at the end of that SST, it follows that slewing to and from the black body (K and K^{-1} slews in Fig. 12), which are unique to the blackbody sequence, do not generate the errors. Therefore, the error (if it occurs) is in A, B or C. Slew B does not often occur (it is related to tracking the limb), thus, the error has to occur in either A or C. Slews A and C are a very often occurring pair, occurring ~ 20 times more frequently than the blackbody pair. Approximately 70% of the total errors occurred in the A-C pair immediately after a black body scan, even though this combined pattern occurs in only 5% of the total slewing. Equally interesting is that 95% of the cases where the pointing error is known, it has the same sign.

The cause of this preconditioning is not understood. Initially, 20% of the black body sequences preconditioned an error. When the errors were the most frequent, all black body sequences

preconditioned a position error. Extensive analysis of the MCS software, electronics and mechanical assembly, has failed to point out any reason for such a preconditioning to occur.

Current Status and Conclusions

Currently the instrument is producing ~90% of the expected science (including all of the atmospheric science). Thus, it was decided to forego further experimentation with the flight instrument. It was determined that the next testing would significantly risk the current actuator performance and would be unlikely to shed any light on the cause or behavior of the anomaly. Instead, it was decided to run the instrument with some operational restrictions to gain as much science as possible through the end of the primary science phase (Dec. 2008). There are two operational restrictions: The first is to use the new SST tables developed during the investigation. The second is to limit scanning in elevation to no more than 10° below the limb. This has prevented further errors.

There is currently no identified root cause for the position errors. All conceived investigations with EM hardware, life test actuator, or in-flight evaluations have failed to shed further light on the issue. Likewise, design reviews and theoretical studies have been exhausted without identifying an obvious root cause.




Figure 12: Slews Implicated in a Typical Position Error

While four potential root causes have not been completely eliminated, none appear to be the answer. None of them matches all observed or anticipated behavior of the hardware. All have major issues that would be sufficient to eliminate them from consideration.

It also should be noted that none of the potential error sources could explain the probable overshoots. It is very difficult to generate sufficient un-commanded movement in the motor to create a 4-step overshoot. While the motor controller could do so, the same software, firmware, and hardware are acting as motor controller for both actuators. Thus if it generates overshoots, it would do so for both actuators. Detailed independent reviews and extensive testing of the motor controller has not shown any propensity to generate overshoots.

The remaining potential root causes are:

- *Keystoning*: It has been demonstrated that keystoning produces position dependent behaviors. But it only occurs at the output bearings (where the available torque is high) and could be expected to affect both actuators (as well as the life test actuator). Keystoning, if it manages to create position errors, would most likely produce large errors (stalling), not 4-step errors, and would never create overshoots.

- *Excessive Friction/Debris*: Excessive mechanical Impedance in transmission or driven load, specific to the flight elevation system, is a possibility. Its origin might be foreign particles/debris in the actuator, parts rubbing each other due to positional shifts, a somewhat loose thermal blanket, workmanship issues, etc. It could explain some positional dependence and subsequent disappearance of the errors with the higher current tables now in place. However, it is difficult to imagine how this failure mode could generate the short four-step nature of most errors, nor would it produce overshoots.

- *Drive Electronics Parts Intermittent Failures*: It is possible that the errors are due to a (partial) failure of an electronics part of the elevation motor driver. This might be due to the solar flare. But the motor driver is completely position insensitive. Besides, none of the known failure modes of any electronics part used has a history of intermittent behavior, nor would it self-correct over time.

- *Controller-Motor Interactions*: The errors could be caused by a subtle interplay between the waveforms generated by the motor controller software, the interpretation of them by the motor drive electronics and the momentum, motion, and normal phase lag of the rotor. Testing and reviews were unable to identify any specific categories of interactions (e.g. resonances) that might produce this behavior, but due to the complex and poorly modeled interfaces involved, there might be an interaction that was not noticed or did not exist with the test equipment, that is slightly dissimilar to the flight hardware. This type of phenomena might be expected to affect both actuators as well. This scenario has no obvious mechanism for generating position dependent errors either. It is also not obvious why it would suddenly appear after several months of perfect operations, get significantly worse, then disappear again.

References

[1] <http://mars.jpl.nasa.gov/mro/>

[2] http://www.planetary.org/programs/projects/mars_exploration/mars_climate_sounder/

[3] http://www.planetary.org/programs/projects/mars_climate_sounder/objectives.html

[4] McCleese, D. J., J. T. Schofield, F. W. Taylor, S. B. Calcutt, M. C. Foote, D. M. Kass, C. B. Leovy, D. A. Paige, P. L. Read, and R. W. Zurek (2007), Mars Climate Sounder: An investigation of thermal and water vapor structure, dust and condensate distributions in the atmosphere, and energy balance of the polar regions. J. Geophys. Res., 112, E05S06, doi:10.1029/2006JE002790.

[5] Pennzane is a product of Nye Lubricants Inc., Fairhaven, MA.

[6] Jau, B. M., "Mechanical Description of the Mars Climate Sounder Instrument." IEEE Aerospace Conference, Big Sky, Montana, March 1-8, 2008.

[7] The Joint Drives were supplied by Starsys Inc., Louisville, CO.

[8] Bearings supplied by Timken, Keene, NH.

Picture Credit

Figures 6 and 7: Courtesy of Starsys Inc.

Acknowledgement

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Copyright 2008 California Institute of Technology.